Minutes of the Database and Related Utilities meeting 11/15/02:

Attending

Joanne Bogart, Toby Burnett, Dave Davis, Traudl Hansl-Kozanecka, Pat Nolan, Bob Schaefer, Karl Young

Minutes:

Item 1: Review of action items

Bob

Checked into using LHEA's Bewoulf to continue prototyping the level 1 database. Former Beowulf administrator is too busy to help us. I am trying to find out if anyone from that group can help us get administrative support for using that cluster. I am not optimistic that that will get much help from them. I think we will make faster progress with a few machines and testing things ourselves. We are likely probably able to get some trashed machines that we can use for testing.

MPI vs. PVM. MPI vs. PVM web page will be posted by Wednesday next week which will describe my findings on PVM and MPI. Overall my impression was that PVM is more tolerant of faults (disk failures, node crashes) and that will be important to our use of a cluster. PVM is also more tolerant of mixed architecture machines – feature which may be important as we will likely add nodes of slightly different architecture as time goes on. PVM is still alive and well.

This provoked some comments from Karl:
He was worried that PVM is too lightweight for our purpose and recommended further study of the issue.

I have queried the leaders of the other working groups asking them to determine What queries they want to make of the Level 1 event database and what catalogs they intend to use. I will compile the responses I get before the next meeting.

I asked HEASARC what protocol they use for making remote database queries. Their standard is to send arguments to a cgi script which then talks to the database. Any other method is liable to be blocked by changing security regulations.

Discussion:

Bob: Karl is it really a problem to do things this way at SLAC?

Karl: It may be able to be done, but there are stringent security regulations that must be followed. I would be required to take a security class in order to be allowed to use it.

Karl:

Checked into identifying some machines to make a mini-cluster at SLAC. May be able to get some limited help from Alex Schlesinger.

Looked up old notes on SCSI disks – they are now irrelevant because the notes are on the SCSI – 2 standard while the current standard is SCSI-3.

Comment from Bob:

I was interested in collecting current level of information so we have a baseline of knowledge on the topic. I am not sure it is worth spending too much time researching it now. As we said last week, t he hardware will likely change by the time we buy it.

Has been doing some research on cluster management software – will post a web page by Wednesday about it.

Dave:

Has talked to SSC about getting test hardware. This is no support for buyiung test machines – other solutions are taking existing machines and tying them together as a cluster and getting machines from excess (trash) to play with.

David Band, science lead for the SSC has been trying to get the name of the person at ASI for just a general contact. He has been ignored so far as they are very busy with SWIFT and are not ready to think about GLAST yet.

Pat:

Has collected links about disks on his Wiki page. The storagereview.com link is a fairly comprehensive site that gives a good overview of the hardware,

Item 2: Meeting time

The meeting time will be shifted so Traudl doesn't have to participate at such a late hour. Furthermore, in two weeks, the Friday occurs right after Thanksgiving, so we will have the next meeting at 11:30 EST on Friday Dec. 6.

Item 3: Response to Software Review

Bob went over the web page summarizing the concerns of the committee and a draft response to those concerns.

Comments from Karl:

The output to ROOT format should be fairly trivial (something Heather could knock off in an afternoon).

We have discussed average number of requests per day, but referred to the fact that we haven't specified how we will deal with requests bunched in time. The higher level of control available in MPI should help us deal with these cases more easily.

Comments from Dave:

As we have used the Chandra data requests as our basis for minimal requirements, I have noticed that it can take hours to retrieve observational data from Chandra. This time is mainly due to the time spent servicing the query and staging the resulting data; the ftp of the data only takes a few minutes. Therefore, our requirements numbers, which are typically better than a few hours should be completely satisfactory as they stand.

Discussion:

Karl: In fact I would like to make a diagram of how we handle the various cases seen by the database.

Bob: I would too – I have been worrying about handling hardware faults. Maybe we can both circulate our ideas to the glastdb mailing list and stimulate some discussion.

Item 4. D2 Livetime, pointing database

D2 Contents:

Pat: I have reviewed the contents of the database in the context of re-examining the likelihood tool and found that everything we need is there or can be derived from what is there. I can't figure out why mission elapsed time is there.

Bob: At least for XTE, mission elapsed is a standard index used all over mission operations and data processing. Perhaps it was left there for other unanticipated uses of D2, which is accessible as a standalone database.

Pat: perhaps we should have UT in the database and not MET?

Bob: I will ask Seth what the motivation for having MET was.

D2 Simulation:

Toby Burnett's mike didn't work, but he has done work simulating the D2 data up and including loading an SQL database. He will send us a report by e-mail about what he has done.

D2 prototyping:

Bob: Dave Davis and I are interesting in prototyping D2 with MySQL. We are interesting in learning about the features and possible imi9tations of MySQL for other reasons so we would be happy to do it.

Karl: I would like to be involved with this prototyping.

Item 5: DB web pages.

Bob: Are web pages adequate?

Traudl: We have no ability to link presentations to the agenda page the way we do on SLAC web pages. Then we keep the links around for a few months as a working record of our results.

Bob: I can contact Richard about moving our Web pages to SLAC. GSFC is rapidly moving to enforcement of section 508, the requirement that pages must be readable by the disabled. I would rather not take responsibility for verifying that the group pages are section 508 complaint, and failure to follow this standard can lead to having your web site be blocked, so it may be worthwhile moving the pages to SLAC. I will contact Richard about this.

Item 6: Report from Traudl re: mirror sites.

Traudl talked to people at Strassbourg about what would be needed for a European mirror site. She has volunteered to be a contact for people who have a lot of experience serving data and will try to learn their recommended techniques. They have asked her to send them a description of our level 1 database. She will draft a 1-2 page document which describes:

- selection criteria (what kind of queries)
- selection machines
- what kind of storage
- how much data we will store and search.

She will find out what access tools are available, what procedures for selecting data are needed.

Discussion:
Bob: Is there any requirements that we use Oracle for our database?
Traudl: No. Typically Oracle and MySQL are used to store and search metadata. We don't yet know exactly where the mirrorsite will be located, so some details are yet to be decided.

Bob: See you all in 3 weeks.

Action Items

All

Send comments about the response to the software review committee to Bob

Bob

Continue to pursue finding hardware for testing (both the LHEA cluster and our own machines).
Post web page about MPI vs. PVM.
Collect responses about requested queries and catalogs.
Continue investigating HEASARC for allowed methods of serving databases from the web.
Circulate strawman database designs
Contact Richard about moving DB web pages to SLAC.

Karl:

Post a web page about cluster management software.
Continue efforts to get hardware for testing.
Circulate strawman database designs.

Dave D.

Continue efforts to get hardware for testing.

Toby Burnett

Send summary of work done on D2 to e-mail distribution list.

Traudl

Send a summary description of our database needs to the glastdblist.